

Constructing Coverability Graphs for Time Basic Petri Nets

Matteo Camilli

Dept. of Computer Science
Università degli Studi di Milano, Italy
camilli@di.unimi.it

Abstract. Time-Basic Petri nets, is a powerful formalism for modeling real-time systems where time constraints are expressed through time functions of marking's time description associated with transition, representing possible firing times. We introduce a technique for coverability analysis based on the building of a finite graph. This technique further exploits the time anonymous concept [5,6], in order to deal with topologically unbounded nets, exploits the concept of a coverage of TA tokens, i.e., a sort of ω anonymous timestamp. Such a coverability analysis technique is able to construct coverability trees/graphs for unbounded Time-Basic Petri net models. The termination of the algorithm is guaranteed as long as, within the input model, tokens growing without limit, can be anonymized. This means that we are able to manage models that do not exhibit Zeno behavior and do not express actions depending on infinite past events. This is actually a reasonable limitation because, generally, real-world examples do not exhibit such a behavior.

Keywords: real-time systems, Time Basic Petri nets, infinite-states systems, reachability problems, coverability analysis

1 Introduction

When analyzing a Petri net, a very common question is whether or not the net is bounded. If it is bounded, the net is theoretically analyzable, and its state space is finite. However the net may be unbounded and classic state space methods generates an infinite number of reachable states from these kind of models. Time Basic (TB) Petri nets [11], as classic Place/Transition nets, may be topologically unbounded. The unboundedness happens whenever there exists a place in the net, where it is possible to accumulate an infinite number of tokens during its execution. Coverability graph algorithms are able to deal with such a models and allow us to decide several important problems: the *boundedness* problem (BP), the *place-boundedness* problem (PBP), the *semi-liveness* problem (SLP) and the *regularity* problem (RP) [13,15]. Anyway, for TB nets, this task is complicated by the time domain. In fact, tokens come along with temporal information and, in general, it is not possible to cluster them into an ω symbol without loosing important information about the system's behavior. However, a

technique able to construct a finite symbolic reachability graph (TRG) relying on a sort of time coverage, was recently introduced [5,6] This technique overcomes the limitations of the existing available analyzers for TB nets, based in turn on a time-bounded inspection of a (possibly infinite) reachability-tree. The time anonymous concept [5,6], introduced by such a technique, allow us to overcome the issue of clustering tokens. In fact, time anonymous timestamps do not carry, for definition, any temporal information. Therefore, an infinite number of TA tokens can be clustered together into a TA^ω symbol without loss of information. The technique, introduced in the current report, gives us a means to deal with topologically unbounded TB net models, where the unboundedness refers to places having an infinite number of TA tokens. Such a limitation is actually reasonable, in practice. In fact, this restricts the analyzable models to systems which do not exhibit *Zeno* behavior and do not express actions depending on “infinite” past events.

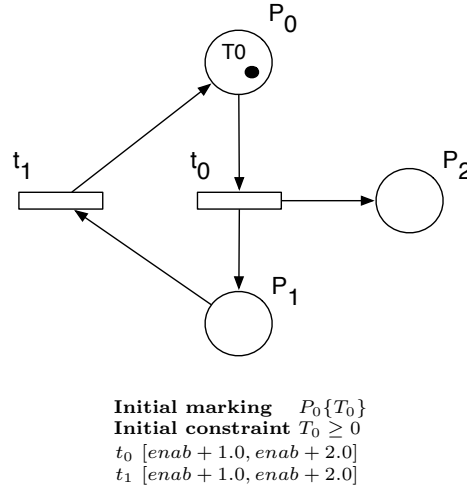


Fig. 1: Simple example showing an unbounded TB net model.

As a simple example, consider the model described in Figure 1. The behavior of the system is very simple: from the initial state, the transition t_0 must fire in the time interval $[T_0, T_0 + 2.0]$. Its firing consumes T_0 and produces two new tokens in places P_1 and P_2 , respectively. In this new state, t_1 is the only enabled transition, and its firing brings the system in the initial topological marking. It is worth noting that every time T_0 fires, a new token is placed into P_2 which cannot be consumed by any firing transition. Therefore, the abstraction technique introduced in [5,6] applied to this example, generates an infinite number of reachable symbolic states because the number of tokens in place P_2 grows without limit. Figure 2 shows a portion of the infinite reachability tree.




Fig. 2: Portion of the infinite reachability tree associated to the TB net model presented in Figure 1.

As we can see, the number of TA tokens in place P_2 grows indefinitely, thus the execution of the software tool GRAPHGEN [5, 6], on such a input, does not terminate. The current report, introduces an extension of the previous analysis technique able to build the coverability graph of unbounded TB nets, exploiting the concept of TA coverage tokens. Our proposal takes inspiration from the Monotone-Pruning (MP) algorithm introduced in [14], for P/T nets, and extends it to deal with TB net models, thus supplying a means, also for real-time systems, to solve the above mentioned problems.

1.1 Preliminaries

A quasi order \geq on a set S is a reflexive and transitive relation on S . Given a quasi order \geq on S , a state $s \in S$ and a subset X of S , we write $s \geq X$ iff there exists an element $s' \in X$ such that $s \geq s'$.

Given a finite set of places P , the marking M ([5, 6]) on P is a function $P \rightarrow Bag(TS \cup \{TA\})$ which supplies foreach place, timestamps associated with tokens. The symbolic ω -marking M^ω on P is a function $P \rightarrow Bag(TS \cup \{TA, TA^\omega\})$. The TA^ω symbol represents, in this case, any number of TA symbols (∞ included). Given the set $U(P) = \mathbb{N}^{|P|}$, an u -marking \bar{u} , is an element of $U(P)$ which associates foreach place, the number of non- TA tokens. Given the set $V(P) = (\mathbb{N} \cup \{\omega\})^{|P|}$, an v -marking \bar{v} , is an element of $V(P)$ which associates foreach place, the number of TA tokens. Given a symbolic state S , we denote with $\bar{u}(S)$, and $\bar{v}(S)$ the u -marking and the v -marking associated with S , respectively.

Given an element $\bar{u} \in U(P)$, $\bar{v} \in V(P)$, and a place p , we denote with \bar{u}_p the number of non- TA tokens in place p , and with \bar{v}_p the number of TA tokens in place p . Since the ω symbol represents an infinite number of TA tokens, the component $\bar{v}_p = \omega$ if and only if $TA^\omega \in M^\omega(p)$.

For instance, if $P = \{p_1, p_2, p_3, p_4\}$ and the symbolic ω -marking is $\{p_1\{T0, TA\}, p_3\{T0, T1, TA^\omega\}\}$, the corresponding u -marking, and v -marking are $\{1, 0, 2, 0\}$, and $\{1, 0, \omega, 0\}$, respectively.

The set $V(P)$ is equipped with a partial order \geq naturally extended by letting $n < \omega, \forall n \in \mathbb{N}$ and $\omega \geq \omega$.

In the current report, when referring to symbolic states, we consider an extended version of the definition proposed in [5, 6], where the marking is represented by the function M^ω rather than M .

Definition 1 (TA erasure). Given a symbolic state $S = \langle M^\omega, C \rangle$, $S_{[-TA]}$ is a symbolic state composed of $\langle M^{\omega'}, C \rangle$, where $M^{\omega'}$ is a symbolic ω -marking obtained from the erasure of all TA symbols from M^ω .

Definition 2 (state coverage). Given two symbolic states $S = \langle M^\omega, C \rangle$ and $S' = \langle M^{\omega'}, C' \rangle$, the u - and v - markings of S \bar{u} , \bar{v} , and the u - and v - markings of S' \bar{u}' , \bar{v}' , S covers S' ($S \geq S'$) iff $\bar{u} = \bar{u}' \wedge \bar{v} \geq \bar{v}' \wedge C \equiv C'$.

That means that S differs from S' only in the number of TA tokens in places. In particular, the number of TA tokens foreach place in S is greater or equal to those ones foreach place in S' . Formally, $\forall p \in P, \bar{v}_p \geq \bar{v}'_p$. Whenever $S \geq S'$ and $\bar{v} \neq \bar{v}'$ we say that S properly covers S' , and we denote it with $S > S'$.

Definition 3 (Coverability Tree). Given a TB net $\mathcal{R} = \langle P, T, F \rangle$, a coverability tree is a tuple $\mathcal{T} = \langle N, n_0, E \rangle$, where N is a set of symbolic states, $n_0 \in N$ is the root state, $E \subseteq N \times T \times N$ is the set of edges labeled with firing transitions. Where:

1. foreach reachable symbolic state S in $TRG(\mathcal{R})$, there exists $S' \in N$ s.t. either $S \subseteq S'$ or $S \geq S'$.
2. foreach symbolic state $S = \langle M^\omega, C \rangle \in N$, having u -marking \bar{u} and v -marking \bar{v} , there exists either a reachable state s of \mathcal{R} s.t. $s \in S$, or a an infinite sequence of reachable symbolic states in $TRG(\mathcal{R})$, $(S_n)_{n \in \mathbb{N}}$ s.t. $\forall n, C_n \equiv C$ and $\forall n, \bar{u}(S_n) = \bar{u}$ and the sequence $(\bar{v}(S_n))_{n \in \mathbb{N}}$ is strictly increasing converging to \bar{v} .

Given a symbolic state $S \in N$, we denote by $Ancestor_{\mathcal{T}}(S)$ the set of ancestors of S in \mathcal{T} (S included). If S is not the root of \mathcal{T} , we denote by $parent_{\mathcal{T}}(S)$ its first ancestor in \mathcal{T} . Finally, given two symbolic states S and S' such that $S \in Ancestor_{\mathcal{T}}(S')$, we denote by $path_{\mathcal{T}}(S, S') \in E^*$ the sequence of edges leading from S to S' in \mathcal{T} .

1.2 Coverability Tree Algorithm

This section presents the algorithm able to construct coverability trees of TB nets. We call it *TBCT* (Algorithm 1) and it is inspired by the Monotonic pruning (MP) algorithm introduced in [14], able to build minimal coverability sets for classic P/T nets. Our proposal involves the acceleration function *Acc*, first introduced in the Karp and Miller algorithm [13]. However, it is defined and also applied in a slightly different manner, in order to deal with a different modeling formalism. In the current context, the *Acc* function, actually modifies the symbolic ω -marking M^ω of a symbolic state by inserting proper TA^ω symbols, accordingly to the following:

$$Acc : 2^N \times N \rightarrow N, Acc(W, S)(p) = S' \text{ s.t.}$$

$$\forall p \in P, \bar{v}(S')_p = \begin{cases} \omega & \text{if } \exists S'' \in W : S'' < S \wedge \bar{v}(S'')_p < \bar{v}(S)_p \wedge S'' \triangleright S \\ m_S(p)_2 & \text{otherwise} \end{cases} \quad (1)$$

Where $S'' \triangleright S$ iff there exists $\sigma = \text{path}_{\mathcal{T}}(S'', S)$, such that σ is feasible from S . Such a condition holds whenever, either:

1. $C_{S''} \implies C_S$, meaning that, $S''_{[-TA]} \subseteq S_{[-TA]}$. In this case, all the paths starting from S'' are feasible from S .
2. $C_S \implies C_{S''}$ and the first component of σ is of type \mathbf{A}^* [4]. In this case $S_{[-TA]} \subseteq S''_{[-TA]}$, therefore not all paths starting from S'' are feasible from S , but since σ starts from all ordinary states of S'' , σ is feasible also from S .

For instance, considering the example in Figure 2, the evaluation of the *Acc* function on $S2$ and its ancestors: $\text{Acc}(\{S0, S1\}, S2)$, causes the insertion of the TA^ω symbol into P_2 because $S2 > S0$, $\bar{v}(S2)_{P_2} > \bar{v}(S0)_{P_2}$ and the path from $S0$ to $S2$ is feasible from $S2$. This way, we recognize that TA tokens into place P_2 can grow without limit.

Algorithm 1 TBCT Algorithm.

Require: A TB net $\mathcal{R} = \langle P, T, F \rangle$

Ensure: A coverability tree $\mathcal{T} = \langle N, n_0, E, L \rangle$, $N = \text{Act} \cup \text{Inact}$

```

1: function TBCT( $\mathcal{R}$ )
2:    $r = \text{BuildRoot}(\mathcal{R})$ 
3:    $N = \{r\}$ ;  $\text{Act} = N$ ;  $\text{Wait} = N$ ;  $E = \emptyset$ ;  $L = \emptyset$ ;
4:   while  $\text{Wait} \neq \emptyset$  do
5:      $s = \text{Pop}(\text{Wait})$ ;
6:     if  $s \in \text{Act}$  then
7:       for  $t \in \text{EnabledTransitions}(s, \mathcal{R})$  do
8:          $m = \text{Successor}(s, t)$ ;
9:          $n = \text{Acc}(\text{Ancestors}_{\mathcal{T}}(m) \cap \text{Act}, m)$ ;
10:         $N+ = \{n\}$ ;  $E+ = \{(s, t, n)\}$ ;
11:        if  $\nexists a \in \text{Act} : a \supseteq n$  then
12:          if  $\exists a \in \text{Act} : a \subset n$  then
13:             $\text{Act}- = \{x : a \in \text{Ancestors}_{\mathcal{T}}(x)\}$ ;
14:          end if
15:          if  $\nexists a \in \text{Act} : a \geq n$  then
16:             $\text{Act}- = \{x : \exists y \in \text{Ancestors}_{\mathcal{T}}(x) \text{ s.t. } y \leq x$ 
17:               $\wedge (y \in \text{Act} \vee y \in \text{Ancestors}_{\mathcal{T}}(n))\}$ ;
18:             $\text{Act}+ = \{n\}$ ;  $\text{Wait}+ = \{n\}$ 
19:          end if
20:        end if
21:      end for
22:    end if
23:  end while
24: end function

```

Likewise both the Karp and Miller and the MP Algorithms, the *TBCT* algorithm builds a coverability tree, but nodes, in the current context, are symbolic states containing symbolic ω -markings and edges are labeled by transitions of the analyzed TB net. Therefore it proceeds by exploring the reachability tree

of the net [5, 6], and accelerating along branches to reach “limit” symbolic ω -markings (containing proper TA^ω symbols). In addition, during the exploration, it can prune branches that are covered by nodes on other branches. Therefore, nodes of the tree are partitioned in two subsets: *active* nodes, and *inactive* ones. Intuitively, active nodes will form the coverability set of the TB net, while inactive ones are not part of the final coverability set, because they are dominated by other active nodes.

The Algorithm 1 proceeds in the following steps to decide how to change the structure \mathcal{T} according to new computed reachable symbolic states:

1. The symbolic state s , popped from *Wait* should be active (test of Line 6).
2. The algorithm iterates through all the enabled transitions and computes one by one all the successor symbolic states: $m = \text{Successor}(s, t)$; (Line 8).
3. The state m is accelerated w.r.t. its active ancestors. A new symbolic state n is created by this operation: $n = \text{Acc}(\text{Ancestors}_{\mathcal{T}}(m) \cap \text{Act}, m)$; (Line 9).
4. If the new symbolic state n is not included or equal to one of the existing active nodes, then it is candidate to be active (test of Line 11).
5. If the new symbolic state n includes an existing active node a , then the sub-tree with root a is deactivated (Lines 12-13).
6. The new symbolic state n is declared as active iff it is not covered by any existing active nodes (test of Line 15 and Line 18).
7. If n is not covered, some symbolic states are deactivated (Line 17).

The update of the set *Act*, complies with the following rules. Intuitively, nodes (and their descendants) are deactivated if they are included or covered by other nodes. This would lead to deactivate a node x iff it owns an ancestor y dominated by n , i.e. such that either $y \subset n$ (Lines 12- 13) or $y \leq n$ (15-17). Concerning the latter case, whenever a new node n (obtained from *Wait*) covers a node y ($y \leq n$), then y can be used to deactivate nodes in two ways:

- if $y \notin \text{Ancestors}_{\mathcal{T}}(n)$, then no matter whether y is active or not, all its descendants are deactivated (Figure 3a).
- if $y \in \text{Ancestors}_{\mathcal{T}}(n)$, then y must be active ($y \in \text{Act}$), and in that case all its descendants are deactivated, except node n itself as it is added to *Act* (Line 18). We require $y \in \text{Act}$ to avoid useless operations. In fact, descendants of n dominates descendants of y (Figure 3b).

For example, considering the example in Figure 2, the insertion of $S2$ accelerated causes the deactivation of both $S0$ and $S1$ because of the execution of line 17. In particular, such a situation corresponds to Figure 3b, because $S2 \geq S0$ and $S0$ (*active* node) belongs to $\text{Ancestors}_{\mathcal{T}}(S2)$.

Figure 4 depicts the coverability tree constructed from the TB net example presented in Figure 1. Elliptic symbolic states form the final coverability set (*active* nodes), while the squared ones are symbolic states deactivated during the analysis. As we can see, the *TBCT* algorithm builds a finite tree structure from an unbounded TB net model. In particular, as shown before, the algorithm is able to identify that the number of TA tokens in place P_2 can grow without limit.

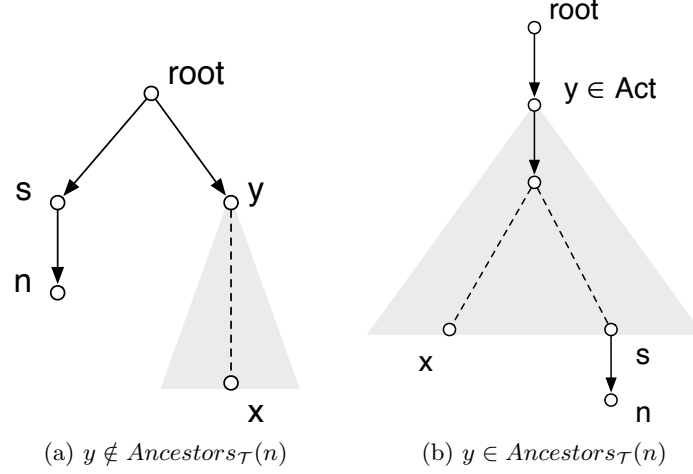


Fig. 3: Deactivations of symbolic states in the TBCT Algorithm




Fig. 4: Coverability tree constructed from the TB net example presented in Figure 1.

As we can see in Figure 4, edges carry information about their type (either AA, EE, AE or EA [4]), and about the local minimum-maximum firing time. In the following, given an edge e , we refer to these information with $\text{type}(e)$ and $\text{time}(e)$, respectively. In particular we refer to the source type with $\text{type}(e)_{\text{src}}$ and to the target type with $\text{type}(e)_{\text{tgt}}$.

It is also possible to construct a coverability graph \mathcal{G} rather than a tree. This task, starting from the tree structure $\mathcal{T} = \langle N, n_0, E \rangle$, executes the following steps:

1. All *inactive* nodes are erased from N .
2. $\forall a \in \text{Act}, \forall \langle a, t, b \rangle \in E$, if b is *inactive*, we search for $a' \in \text{Act}$ so that $a' \geq b$ or $a' \geq b$, thus we remove $\langle a, t, b \rangle$ from E and we insert $\langle a, t, a' \rangle$.
3. All covered edges (Definition 4) are removed from E .

Definition 4 (edge coverage). Given a coverability tree $\mathcal{T} = \langle N, n_0, E \rangle$ and two edges $e = \langle a, t, b \rangle$, $e' = \langle a', t', b' \rangle \in E$, e covers (\geq) e' iff:

- i $a = a' \wedge t = t' \wedge b = b'$
- ii $\text{time}(e) \geq \text{time}(e')$

iii $\text{type}(e)_{src} \geq \text{type}(e')_{src} \wedge \text{type}(e)_{trgt} \geq \text{type}(e')_{trgt}$, being $A > E$




Fig. 5: Coverability graph constructed from the coverability tree presented in Figure 4.

Figure 5 shows the coverability graph resulting from the coverability tree presented in figure 4. Such a structure contains only active symbolic states and gives us a more intuitive overview on the system's behavior. For instance, by observing Figure 5, it's easy to figure out that the system alternates two symbolic states where P_0 and P_1 are marked with a single token, while place P_2 can accumulate TA tokens without limit.

The rest of this section reports some simple examples of unbounded TB net models analyzed by the software tool implementing the *TBCT* algorithm. All the coverability trees/graphs have been automatically obtained by using *GRAPHVIZ* visualization software [12] on the output generated from the tool-set. The *TW* notation used into symbolic ω -markings, stands for TA^ω .

Example A Figure 6 depicts an unbounded TB net model with two places (P_0 , P_1) and two transitions (T_0 , T_1). It represents a simple synchronous system, where an operation occurs at each time unit (e.g. production/consumption). Produced units are stored into a infinite buffer. After the first consumption the system stops.

Figure 7a shows the coverability tree of *A*. As we can see, the introduction of S_1 causes the deactivation of S_0 ($S_1 \geq S_0$). From S_1 the system can evolve either into S_2 which is inactive ($S_2 = S_1$), or S_3 which is a final state. Such a behavior is also shown by the coverability graph (Figure 7b): the system loops into S_1 , by the firing of T_0 transition, until the firing of T_1 which leads into the final state S_3 .

Example B This model (Figure 8) is analogue to model *A*, except for an additional arc and a different initial marking. It represents two synchronous tasks, where each task can either produce or consume. An infinite buffer stores produced units. Figure 9a and 9b show its coverability tree and coverability graph, respectively. It is worth noting that the firing of T_0 from S_1 produces an additional token into place P_1 and because of the recognition of both tokens of




Fig. 6: Unbounded TB net model A.

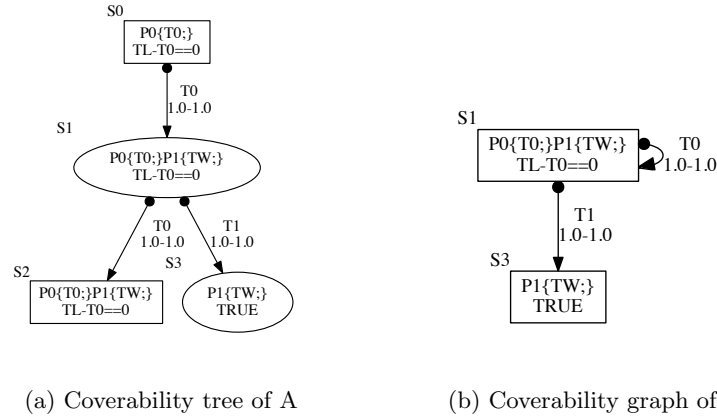



Fig. 7: Coverability tree/graph of example A.

$P1$ as TA , the acceleration of $S2$ recognizes the TA^ω into $P1$. Therefore, $S2$ deactivates both $S0$ and $S1$. Successors of both $S3$ and $S4$ are identified equal to $S2$.

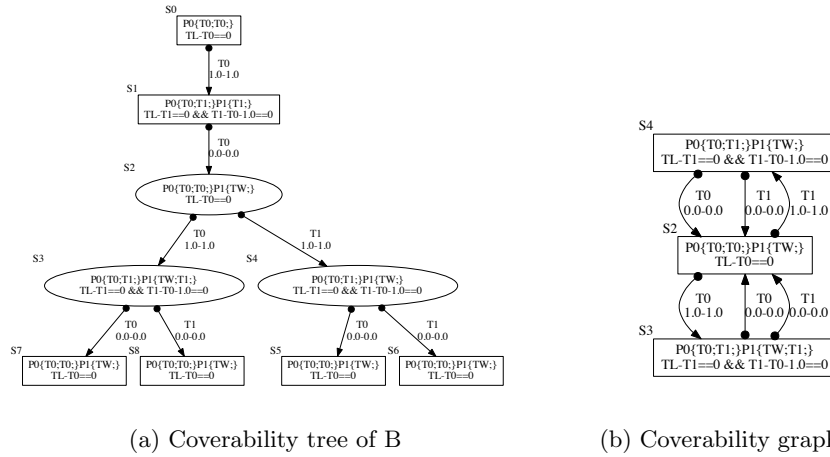
Example C This model (Figure 10) represents an unbounded TB net with four places ($P0, P1, P2, P3$), two strong transitions ($T0, T2$) and a weak transition ($T3$). Transition $T0$ acts as a sort of timer, in fact, whenever enabled, it must fire in 10 time units from its previous firing time. Figure 11a and 11b show its coverability tree and coverability graph, respectively.

Concerning the current example, it is worth noting that before the introduction of $S4$, all the symbolic states were *active*. The acceleration of $S4$ leads to the recognition of a TA^ω into place $P3$, and thus the identification of the coverage $S4 \geq S1$. This causes the deactivation of both $S1$ and its descendants $S2$ and $S3$. The successors of $S4$ are $S5$ and $S6$. In this case, since $S5 \subset S6$, $S5$ is deactivated. Finally, $S7$ is recognized to be equal to $S4$.



Initial marking $P0\{T0, T0\}$
Initial constraint $T0 \geq 0$
 $T0 [enab + 1.0, enab + 1.0]$
 $T1 [enab + 1.0, enab + 1.0]$

Fig. 8: Unbounded TB net model B.




(a) Coverability tree of B

(b) Coverability graph of B

Fig. 9: Coverability tree/graph of example B.

1.3 Related Work

Concerning the reachability analysis of classic P/T nets, Karp and Miller (K&M) introduced an algorithm for computing the *minimal coverability set* (MCS) [13]. This algorithm builds a finite tree representation of the (potentially infinite) reachability graph of the given P/T net. It uses acceleration techniques to collapse branches of the tree and ensure termination. The K&M Algorithm has been also extended to other classes of well-structured transition systems [8, 9]. Anyway, the K&M Algorithm is not efficient in analyzing real-world examples and it often does not terminate in reasonable time. One reason is that in many cases it will compute several times a same subtree. The MCT algorithm [7] introduces clever optimizations: a new node is added to the tree only if its marking is not smaller than the marking of an existing node. Then, the tree is pruned: each node labelled with a marking that is smaller than the marking of the new



Initial marking $P0\{T0\}P1\{T0\}$
Initial constraint $T0 \geq 0$
 $T0 [enab, P0 + 10.0]$
 $T1 [enab + 2.0, enab + 3.0]$
 $T2 [enab + 1.0, enab + 4.0]$

Fig. 10: Unbounded TB net model C.

node is removed together with all its successors. The idea is that a node that is not added or that is removed from the tree should be covered by the new node or one of its successors. However, the MCT algorithm is flawed [10]: it computes an incomplete *forward* reachability set (i.e. all the markings reachable from the initial markings). In [10], the COVERPROC algorithm, is proposed for the computation of the MCS of a Petri net. This algorithm follows a different approach and is not based on the K&M Algorithm. In [14], the MP algorithm is proposed. This algorithm can be viewed as the MCT algorithm with a slightly more aggressive pruning strategy. Experimental results show that the MP algorithm is a strong improvement over both the K&M and the COVERPROC algorithms. The *TBCT* algorithm, introduced in the current report, is somehow inspired by the MP algorithm, and is able to construct coverability graphs of real-time systems modeled with TB nets.

For timed Petri nets (TPNs), although the set of *backward* reachable states (i.e. all the markings from which a final marking is reachable) is computable [2], the set of *forward* reachable states is in general not computable. Therefore any procedure for performing forward reachability analysis on TPNs is incomplete. In [1], an abstraction of the set of reachable markings of TPNs is proposed. It introduces a symbolic representation for downward closed sets, so called region generators (i.e. the union of an infinite number of regions [3]). Anyway, the termination of the forward analysis by means of this abstraction is not guaranteed.

In the current report, we addressed unbounded TB nets, which represent a much more expressive formalism for real-time systems than TPNs (interval bounds in TB nets are linear functions of timestamps in the enabling marking, rather than simple numerical constants). Other coverability analysis techniques for such a formalism, have not been proposed yet, as far as we know.

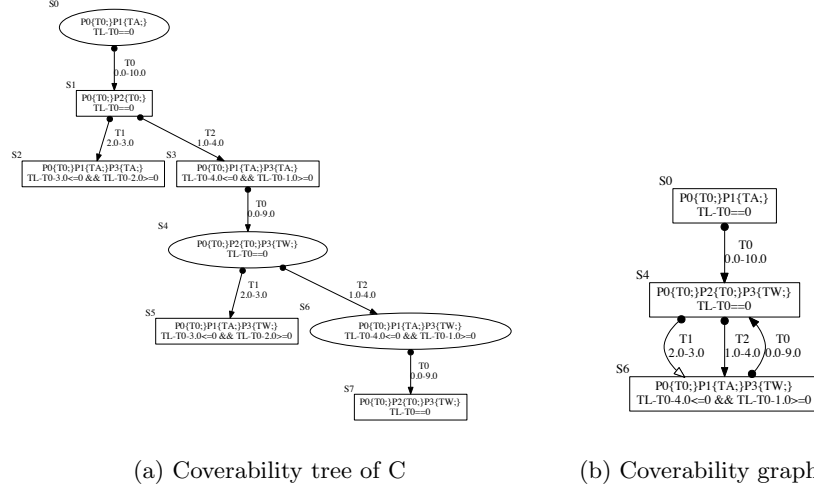


Fig. 11: Coverability tree/graph of example C.

1.4 Conclusion

The current report introduces a coverability analysis technique able to construct a coverability tree/graph for unbounded TB net models. The termination of the *TBCT* algorithm is guaranteed as long as, within the input model, tokens growing without limit, can be anonymized. This means that we are able to manage models that do not exhibit *Zeno* behavior and do not express temporal functions depending on “infinite” past events. This is actually a reasonable limitation because, in general, real-world examples do not exhibit such a behavior.

References

1. Parosh Aziz Abdulla, Johann Deneux, Pritha Mahata, and Aletta Nylén. Using forward reachability analysis for verification of timed petri nets. *Nordic J. of Computing*, 14(1):1–42, January 2007.
2. Parosh Aziz Abdulla and Aletta Nylén. Timed petri nets and bqos. In *Proceedings of the 22Nd International Conference on Application and Theory of Petri Nets, ICATPN '01*, pages 53–70, London, UK, UK, 2001. Springer-Verlag.
3. Rajeev Alur and D. L. Dill. Automata for modeling real-time systems. In *Proceedings of the Seventeenth International Colloquium on Automata, Languages and Programming*, pages 322–335, New York, NY, USA, 1990. Springer-Verlag New York, Inc.
4. Carlo Bellettini, Matteo Camilli, Lorenzo Capra, and Mattia Monga. Mardigras: Simplified building of reachability graphs on large clusters. In Parosh Aziz Abdulla and Igor Potapov, editors, *Reachability Problems*, volume 8169 of *LNCS*, pages 83–95. Springer Berlin Heidelberg, 2013.

5. Carlo Bellettini and Lorenzo Capra. Reachability analysis of time basic petri nets: A time coverage approach. In *Proceedings of the 2011 13th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, SYNASC '11*, pages 110–117, Washington, DC, USA, 2011. IEEE Computer Society.
6. M. Camilli. Verification of Reachability Problems for Time Basic Petri Nets. *ArXiv e-prints*, September 2014.
7. Alain Finkel. The minimal coverability graph for petri nets. In *Papers from the 12th International Conference on Applications and Theory of Petri Nets: Advances in Petri Nets 1993*, pages 210–243, London, UK, UK, 1993. Springer-Verlag.
8. Alain Finkel and Jean Goubault-Larrecq. Forward analysis for wsts, part i: Completions. In Susanne Albers and Jean-Yves Marion, editors, *26th International Symposium on Theoretical Aspects of Computer Science, STACS 2009, February 26-28, 2009, Freiburg, Germany, Proceedings*, volume 09001 of *Dagstuhl Seminar Proceedings*, pages 433–444. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2009.
9. Alain Finkel and Jean Goubault-Larrecq. Forward analysis for wsts, part ii: Complete wsts. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming: Part II, ICALP '09*, pages 188–199, Berlin, Heidelberg, 2009. Springer-Verlag.
10. Gilles Geeraerts, Jean-François Raskin, and Laurent Van Begin. On the efficient computation of the minimal coverability set for petri nets. In *Proceedings of the 5th International Conference on Automated Technology for Verification and Analysis, ATVA'07*, pages 98–113, Berlin, Heidelberg, 2007. Springer-Verlag.
11. Carlo Ghezzi, Dino Mandrioli, Sandro Morasca, and Mauro Pezzè. A unified high-level petri net formalism for time-critical systems. *IEEE Trans. Softw. Eng.*, 17:160–172, February 1991.
12. <http://www.graphviz.org/>. Graphviz - graph visualization software.
13. Richard M. Karp and Raymond E. Miller. Parallel program schemata. *J. Comput. Syst. Sci.*, 3(2):147–195, May 1969.
14. Pierre-Alain Reynier and Frédéric Servais. Minimal coverability set for petri nets: Karp and miller algorithm with pruning. In *Proceedings of the 32Nd International Conference on Applications and Theory of Petri Nets, PETRI NETS'11*, pages 69–88, Berlin, Heidelberg, 2011. Springer-Verlag.
15. Rüdiger Valk and Guy Vidal-Naquet. Petri nets and regular languages. *Journal of Computer and System Sciences*, 23(3):299 – 325, 1981.